What is claimed is:

1. A method of communicating between nodes in a clustered computer system, the method comprising:

(a) communicating a port identifier from a first node to a second node coupled to the first node over a point-to-point network, wherein the first node includes a plurality of network ports and a plurality of communication registers, wherein each communication register is dedicated to an associated network port among the plurality of network ports and is configured to store data received over such associated network port, and wherein the port identifier identifies a network port among the plurality of network ports to which the second node is coupled to the first node; and

(b) communicating data from the second node to the first node by initiating a write operation on the first node using the second node to store the data in the communication register associated with the network port identified by the port identifier.

2. The method of claim 1, further comprising detecting in the first node the storage of data in the communication register associated with the network port identified by the port identifier.

3. The method of claim 2, further comprising generating an interrupt on the first node in response to detecting the storage of data in the communication register associated with the network port identified by the port identifier.

4. The method of claim 3, further comprising processing the interrupt by processing the data stored in the communication register associated with the network port identified by the port identifier, and clearing the interrupt.

5. The method of claim 4, wherein detecting the storage of data comprises detecting a non-zero value stored in any of the plurality of communication registers,

and wherein clearing the interrupt comprises resetting the plurality of communication registers to zero values.

6. The method of claim 1, wherein communicating the data comprises sequentially storing a plurality of commands in the communication register associated with the network port identified by the port identifier, the method further comprising processing each of the plurality of commands in the first node.

7. The method of claim 1, further comprising initiating, with the second node, a read operation for a configuration register in the first node, wherein communicating the node identifier is performed in response to the read operation.

8. The method of claim 1, wherein communicating the node identifier is performed in response to a read request sent over the point-to-point network by the second node.

9. The method of claim 1, wherein the plurality of communication registers are allocated a range of register addresses in a register address space for the node, and wherein communicating the data comprises sending a write request to the register address of the communication register associated with the network port identified by the port identifier.

10. A circuit arrangement, comprising:

(a) a plurality of network ports, each configured to couple a first node from a clustered computer system to another node in the clustered computer system over a point-to-point network;

(b) a plurality of communication registers, each dedicated to an associated network port among the plurality of network ports and configured to store data received through such associated network port; and

(c) a control circuit coupled to the plurality of communication registers and configured to automatically notify the first node in response to storage of data in any of the plurality of communication registers.

11. The circuit arrangement of claim 10, wherein the control circuit is configured to detect the storage of data in a communication register among the plurality of communication registers by detecting a non-zero value stored in such communication register.

12. The circuit arrangement of claim 11, wherein the control circuit is configured to automatically notify the first node by generating an interrupt.

13. The circuit arrangement of claim 12, wherein the control circuit is configured to generate a common interrupt for all of the plurality of communication registers.

14. The circuit arrangement of claim 10, wherein each communication register includes a plurality of binary outputs, and wherein the control circuit comprises at least one logic gate configured to generate an interrupt signal by performing a logical-OR operation on all of the binary outputs of the plurality of communication registers.

15. The circuit arrangement of claim 10, wherein the control circuit is further configured to output a port identifier over a first network port among the plurality of network ports in response to a read request received over the first network port, the

port identifier identifying the first network port as the network port from which the read request was received.

16. The circuit arrangement of claim 15, further comprising a configuration register, wherein the control circuit is configured to output data stored in the configuration register in response to the read request.

17. The circuit arrangement of claim 16, wherein the plurality of communication registers are allocated a range of register addresses in a register address space for the node, and wherein the control circuit is configured to store data received over a first network port among the plurality of network ports in the communication register associated with the first network port in response to a write request addressed to the register address of the communication register associated with the network port identified by the port identifier.

18. An integrated circuit device comprising the circuit arrangement of claim 10.

19. A program product comprising a hardware definition program that defines the circuit arrangement of claim 10, and a signal bearing medium bearing the hardware definition program, wherein the signal bearing medium includes at least one of a transmission medium and a recordable medium.

20. A node for use in a clustered computer system, the node comprising:

(a) a plurality of network ports, each configured to couple to another node in the clustered computer system over a point-to-point network;

(b) a plurality of communication registers, each dedicated to an associated network port among the plurality of network ports and configured to store data received through such associated network port; and

(c) a control circuit coupled to the plurality of communication registers and configured to automatically notify the node in response to storage of data in any of the plurality of communication registers.

21. The node of claim 20, wherein the control circuit is configured to generate the notification by signaling an interrupt in response to any of the plurality of communication registers storing a non-zero value.

22. A clustered computer system comprising:

(a) a plurality of nodes, each node including:

(i) a plurality of network ports;

(ii) a plurality of communication registers, each dedicated to an associated network port among the plurality of network ports and configured to store data received through such associated network port; and

(iii) a control circuit coupled to the plurality of communication registers and configured to automatically notify such node in response to storage of data in any of the plurality of communication registers; and

(b) a plurality of point-to-point network interconnects, each coupled between a pair of nodes from the plurality of nodes through network ports on each of the pair of nodes.